INTERNATIONAL JOURNAL OF COGNITIVE ERGONOMICS, 2001, *5*(3), 317–333

What Do Sharable Instructional Objects Have to do With Intelligent Tutoring Systems, and Vice Versa?

J. D. Fletcher
Institute for Defense Analyses

ABSTRACT

By reviewing data on classroom and tutorial instruction, this article presents a perspective on the value of technology-based instruction in general and intelligent tutoring systems (ITS) in particular. It finds mixed-initiative dialogue and on-demand, real-time generative capabilities to be defining functionalities of ITS. These functionalities are motivated by basic research into human learning and cognition. From this perspective, development of both the Advanced Distributed Learning (ADL) initiative and ITS are viewed as parallel but presently independent activities. Their common interests in the development and availability on the World Wide Web of sharable (accessible, interoperable, durable, and reusable) instructional objects that can act as either instructional content or algorithmic agents are described and discussed. Sharable instructional objects are likely to reduce the costs and increase the effectiveness of both ADL technologies and ITS. Development of these objects should be cooperatively promoted and pursued by both the ADL and ITS communities. Doing so may lead to the development and wide use of personal learning associates, allowing high-quality instruction and decision aiding to become ubiquitous and affordable.

1. INTRODUCTION

The techniques of object-oriented programming are influencing many areas of software development. Development of tools, materials, and applications for technology-based education and training are among these. This article discusses the value of technology-based instruction in general, intelligent tutoring systems (ITS) in particular and how both may be enhanced by the specification and development of sharable instructional objects.

2. WHAT ARE THE CONTRIBUTIONS OF TECHNOLOGY-BASED INSTRUCTION?

Substantial improvements in instructional effectiveness may be obtained by tailoring instruction to the needs and capabilities of individual learners. Evidence for these improve-

Requests for reprints should be sent to J. D. Fletcher, Institute for Defense Analyses, 4850 Mark Canter Drive, Alexandria, VA 22311. E-mail: fletcher@ida.org

ments may be found in studies of one-on-one tutoring. One widely cited discussion was based on studies performed by Benjamin Bloom (1984) and his students comparing the achievement of individually tutored students (one instructor for each student) with that of classroom students (one instructor for every 28 to 32 students). It is not surprising to find that individual tutoring in these studies increased the achievement of students. What is surprising is the magnitude of the increase. Bloom reported that the overall difference in achievement across three studies was about two standard deviations, which means, roughly, that tutoring improved the achievement of 50th percentile students to that of 98th percentile students. Two standard deviations is a large difference. Bloom posed it as a challenge to researchers and developers in education.

Why is this two-sigma difference (as Bloom [1984] called it) a challenge? Why don't we simply provide one-on-one tutoring for all our students? The answer is straightforward and obvious: We can't afford it. The provision of one instructor for each student is, in most cases, prohibitively expensive. Because of this problem, Michael Scriven (1975) early on described individualized, tutorial instruction as both an instructional imperative and an economic impossibility.

We may now have the means to break out of this dilemma. Fundamental to the value of our rapidly evolving computer technology is its ability to adjust its actions in real time and on demand to the conditions of the moment. The promise of this capability for instruction has not been lost on researchers and developers. The ability to adjust actions responsively to meet the needs of individual students is used by both tutorial instruction and technology-based instruction to solve two significant problems that confront all classroom teachers and that may account for much of the difference in effectiveness between individualized, tutorial instruction and classroom instruction—the problems of individualized pace and interactivity.

2.1. The Challenges of Classroom Instruction: Pace

The variability in the times that different students take to learn and that teachers must accommodate in classrooms can be overwhelming. Despite conscientious efforts to sustain common levels of prior knowledge in classrooms, our current school practices appear to increase these differences by about 1 year for every year students spend in elementary school (Heuston, 1997). For instance, the average spread of academic achievement in third grade is about 3 years. By sixth grade, it increases to about 6 years.

Research has helped confirm the extent of this variability. Early on, Suppes (1964) found the time needed by kindergarten students to learn to build words from letters varied by about 13:1. Ratios for Grade 5 students to master a unit of social studies were found to be 3:1 and 5:1 (Gettinger & White, 1980). In two different studies, the rates of learning observed among hearing impaired and Native American students were found to be 4:1 (Suppes, Fletcher, & Zanotti, 1975, 1976). Based on a range of research findings, Carroll (1970) estimated the overall ratio for elementary school students to be 5:1. Even among highly selected students at a major research university, the times needed by undergraduates to learn a programming language have been seen to vary by 7:1 (A. T. Corbett, personal communication, April, 29, 1998). The latter result helps confirm earlier observations that even though rates of learning may be initially determined by basic ability, the effects of ability are quickly overtaken by those of prior knowledge and experience with the subject matter (Tobias, 1989).

What happens when we allow students to progress at their own rates—when we individualize the pace of learning as tutors do and as we have long been able to do using technology-based instruction? Some studies have compared time to achieve specific instructional objectives in classrooms with time to achieve the same objectives using technology-based instruction. Reviews combining the results of these studies have been reported. Orlansky and String (1979) found that reductions in time to reach instructional objectives averaged about 54% in their review of computer-based instruction used in military training. Fletcher (1997) reported an average time reduction of 31% in six studies of interactive multimedia instruction applied in higher education. Kulik (1994) found time reductions of 34% in 17 studies of computer-based instruction used in higher education, and 24% in 15 studies of adult education. All these reviews are effectively independent in that they reviewed different sets of evaluation studies. On this basis, reductions of about 30% in the time it takes students to reach a variety of given instructional objectives seem to be a reasonable expectation.

We should expect that the individualization of pace allowed by technology-based instruction, including ITS, is likely to produce significant savings in time to learn. Do these savings in time matter? Studies have shown that hundreds of millions of dollars might be saved in Department of Defense specialized skill training if the training time to reach specified thresholds of knowledge and skill can be reduced by 30% (National Research Council, 1997). These savings result as much from reductions in training costs as from the value of providing trained personnel ready for duty in operational units earlier.

We might expect such savings to be substantial and operationally significant in military and industrial training where the students are paid by the organization providing the training and where both the students and their employers have a stake in their rapid mastery of subject matter and early availability for organizational operations. But these savings may well be equally valuable in K–12 education. As Gettinger and White (1980) asked, Why should we assume that time to learn is unimportant to schoolchildren? All societies have a heavy stake in their success. Opportunities, such as those provided by technology-based instruction, permitting individual students to expand their capabilities and realize their potential as rapidly as possible, may be even more important than they are in military and industrial training.

2.2. The Challenges of Classroom Instruction: Interactivity

If we consider interactivity to be question-and-answer sequences occurring between students and instructors, might it also help account for the differences we observe between one-on-many classroom instruction and one-on-one tutorial instruction? Those who study classroom interactions of this sort have found that groups of students ask about three questions an hour and that any single student in a class asks about 0.11 questions an hour (Graesser & Person, 1994). By contrast, students in individual tutorial sessions have been found to ask 20 to 30 questions an hour and have been required to answer 117 to 146 questions an hour. Finally, some students completing computer-based instruction may answer 8 to 12 questions a minute—questions that have been especially selected to meet their individual needs and that are immediately graded to provide feedback.

The intensity of tutorial instruction provided by technology evidently pays off in student achievement. In a review of 97 empirical evaluations almost entirely made up of standard (non-ITS) computer-based instruction compared with conventional classroom approaches, Kulik (1994) reported an effect size advantage for computer-based approaches of about 0.39 standard deviations—or roughly an improvement of 50th percentile students to about the 65th percentile of achievement. In an attempt to determine the advantages to instruction added by multimedia capabilities, Fletcher (1997) summarized a review covering 47 comparisons of interactive multimedia instruction with conventional classroom approaches and found an effect size advantage for these technology-based approaches of about 0.50 standard deviations—or roughly an improvement of 50th percentile students to about the 69th percentile of achievement. Fletcher also reported 11 ITS evaluations and found an effect size advantage for ITS of about 0.84 standard deviations—roughly an improvement of 50th percentile students to about the 79th percentile of achievement. In a review covering five evaluations of the SHERLOCK ITS system, Gott, Kane, and Lesgold (1995) reported an overall effect size advantage of about 1.05 standard deviations—roughly an improvement of 50th percentile students to about the 85th percentile of achievement.

We cannot say if we will achieve Bloom's (1984) target improvement of two standard deviations, but the available evidence suggests that we are progressing in the right direction through the use of technology-based instruction. These findings also suggest that ITSs raise the ceiling for the effectiveness of technology-based instruction and provide new avenues for improvement over standard classroom instruction. We may be approaching the limits of improvements we can achieve with general, non-ITS computer-based instruction. We may have much yet to learn about the instructional possibilities for ITS.

2.3. Are We Getting Anywhere? Costs

These results show promise for technology-based instruction in general and for ITS in particular. However, if we seek to make a difference in the day-to-day practice of instruction, we must address the concerns of those who are making decisions about it. For them, effectiveness is an important consideration but only one component of the equation. The hallmark of administrative decision making is consideration of what must be given up to gain some advantage. Generally, these decisions involve trade-offs between costs and effectiveness. Decision making in instruction is no different. Costs must be considered as well as effectiveness.

The ratios of costs for technology-based instruction compared with more conventional instruction can be reported in three categories: initial investment costs to develop and implement both types of training; operating and support costs for both types of training once it is in place; and these two cost categories combined.

In reports structured for this purpose, the smaller the ratio, the better the news for the technology-based training. For one set of studies where achievement under technology-based training was at least equal and generally superior to that produced by more conventional instruction, Fletcher (1997) reported the ratios to be 0.43 for initial investment, 0.16 for operating and support, and 0.35 overall. In these studies, most of the cost savings achieved by technology-based instruction were due to the substitution of simulated equipment for the real equipment to be used on the job. Nonetheless, these ratios suggest that substantial economies can be realized through the use of technology-based instruction and that the return can justify the investment. More complete analyses based on projections suggest that potential savings, not just cost-avoidances, for the U.S. Department of Defense may amount to hundreds of millions of dollars (National Research Council, 1997).

2.4. In Summary: The Thirds

Assessments of technology-based instruction leave us with "the thirds." Use of technology reduces the cost of instruction by about one third, and, additionally, it either reduces time to reach given instructional objectives by one third or it increases the achievement of its students by about one third. The primary payoff for many organizations is, of course, the more rapid preparation of personnel to perform operational duties.

3. WHAT ARE THE CONTRIBUTIONS OF ITS?

It may be best to begin by noting the features that garden-variety computer-based instruction can and does provide. Notably it can (a) accommodate individual students' rate of progress, allowing as much or as little time as each separate student needs to reach instructional objectives; (b) adjust the sequence of instructional content to each student's needs; (c) adjust the content itself—different students can receive different content depending on what they have mastered and what they have yet to learn; (d) make the instruction as easy or as difficult as necessary; and (e) adjust to the learning style (e.g., verbal vs. visual) that is most appropriate for each student. These capabilities have been available and used in computer-based instruction from its inception in the 1950s (e.g., Atkinson & Fletcher, 1972; Fletcher & Rockway, 1986; Suppes & Morningstar, 1972). Those who promote systems with these features, touting them as indicators of newly developed "intelligent" capabilities, may be missing some history. Whatever the case, they are using the term *intelligent* in ways that differ from the historical objectives of ITS—objectives that have been pursued since the late 1960s.

What are these objectives? What is left for ITS to provide? What can we get from it that is not otherwise available? Two functionalities deserve mention:

- The ability to allow either the computer or the student to ask open-ended questions and initiate instructional, "mixed-initiative" dialogue as needed or desired.
- Related to this, the ability to generate instructional material and interactions on demand rather than require developers to foresee and prestore all such materials and interactions needed to meet all possible eventualities.

The first of these functionalities requires the ITS to understand and participate in mixed-initiative interactions with the student. It requires a mutual understanding of a language for information retrieval, decision aiding, and instruction that is shared by both the ITS and the student or user. Natural language has been a frequent choice for this capability, but the language of mathematics, mathematical logic, and electronics have been used, as publications edited by Suppes (1981), Sleeman and Brown (1982), Psotka, Massey, and Mutter (1988), and Farr and Psotka (1992) have reported.

Whatever form it takes, mixed-initiative dialogue in which either the student or the instructor can initiate interactions appears to be a key feature of one-on-one tutorial instruction (e.g., Graesser, Person, & Magliano, 1995). In attempting to secure for our students the benefits of one-on-one tutorials, we may well turn to technology-based instruction that specifically seeks to develop and implement mixed-initiative dialogue as an instructional approach. Such a capability has long been a goal of ITS (Carbonell, 1970).

The second functionality requires ITS to devise on demand—not retrieve from storage—interactions and presentations for individual students. This capability involves more than generating elements to fill in blanks in a template. It means generating interactions and presentations from information primitives using an "instructional grammar" that is analogous to the deep structure grammar of the transformational-generative linguists of a generation ago. This functionality harkens back to the roots of ITS development as (again) can be seen in the volumes edited by Suppes (1981), Sleeman and Brown (1982), Psotka, Massey, and Mutter (1988), and Farr and Psotka (1992).

Motivations for both these functionalities can be found in basic research into human learning, memory, perception, and cognition. Findings from this research have led us to view all cognitive processes as constructive and regenerative. They have caused general theories of perception and learning to evolve from the fairly strict logical positivism of behavioral psychology, which emphasizes the study of directly observable and directly measurable actions, to greater consideration of internal, less observable processes that are assumed to mediate and enable human learning—and to produce the directly observable behavior that is the subject of behaviorist approaches. The keynote of these newer conceptions of cognition may have been struck by Ulric Neisser (1967) who stated, "The central assertion is that seeing, hearing, and remembering are all acts of *construction,* which may make more or less use of stimulus information depending on circumstances" (p. 10).

Neisser (1967) was led to this point of view by a large body of empirical evidence showing that many aspects of human behavior, such as seeing and hearing, simply could not be accounted for by external physical cues reaching human perceptors, such as eyes and ears. Additional processes, including an internally, one might say cognitively, generated analysis by synthesis had to be posited to account for well-established and observable human abilities to detect, identify, and process physical stimuli. Such a process requires an active synthesis of the environment based on a runnable cognitive model, or simulation, that is validated or modified as needed by cues being received from sensory perceptors. It is the actively evolving simulation, not the stimuli alone, that is said to account for what the individual learns.

These ideas were, of course, around long before Neisser (1967) published his book, and it is notable that they did not occur only to psychologists. Norbert Weiner (1954), a cyberneticist, suggested that,

> In both [the living individual and the machine] there exists a special apparatus for collecting information from the outside world at low energy levels, and for making it available for the operation of the individual or of the machine. In both cases these external messages are not taken *neat*, but through the internal transforming powers of the apparatus, whether it be alive or dead. The information is then turned into a new form available for the further stages of performance. (pp. 26–27)

All this leads to the impression that the generative capability sought by ITS and the Advanced Distributive Learning (ADL) initiative is not something merely nice to have, but essential if we are to advance beyond the constraints of the prescribed branching, programmed learning, and the ad hoc principles currently used to design technology-based instruction. A generative approach is essential if we are to deal successfully with the immensity, extent, and variability of human cognition.

The key defining characteristic of both ITS and ADL, then, is not application of computer techniques from artificial intelligence or knowledge representation, or the specification of sharable instructional objects, important as these may be. It is rather the functional capability to generate in real time and on-demand instructional interactions that are tailored to student requests or needs. This generative capability motivated the U.S. Department of Defense to invest in ITS originally (Fletcher & Rockway, 1986). At that time, the motivation was to reduce or eliminate the high costs of foreseeing or predicting all possibly needed materials and interactions, programming them, and prestoring them into computer-based instruction. Today, this motivation remains fundamental to those who support ITS and ADL development.

3.1. An Example of an ITS at Work

An example of these capabilities at work may be found in the student–computer, mixed-initiative dialogues supported by the Sophisticated Instructional Environment (SOPHIE)—an early and seminal ITS (Brown, Burton, & DeKleer, 1982). SOPHIE used a general-purpose electronic simulator to represent an electronic power supply and create a reactive, tutorial environment for training in electronic troubleshooting. The underlying simulator allowed faults introduced into the system by a (human) instructor or by SOPHIE itself to be propagated throughout the (simulated) power supply in a realistic fashion. SOPHIE employed a semantic network of information about the power supply, one of the most sophisticated natural language understanders ever built into ITS, and an articulate, expert troubleshooting capability that could not only solve arbitrary faults introduced by students or instructors but also explain to the student how it solved them.

SOPHIE's reactive environment responded in a realistic fashion to the actions and decisions of a student engaged in practical problem solving. It employed strategies that shadowed those used by the student to make inferences about the circuit based on the student's measurements. SOPHIE could thereby determine which deductions the student was using, construct a model of his or her state of knowledge, and use this model to generate explanations in terms that the student was prepared to understand. SOPHIE highlighted the need for intelligent training systems not only to solve arbitrary problems, but to do so in a way that could be explained to students and understood by them. It pioneered capabilities currently lacking in many information acquisition and retrieval systems.

The following is a student-computer dialogue adapted from SOPHIE:

Student: What is the base emitter voltage of the voltage limiting transistor?
SOPHIE: The base emitter voltage of Q5 is –.037 volts.

The first student–SOPHIE interaction in this example is a straightforward question and answer. SOPHIE has introduced a fault into the simulated power supply, and the student is trying to locate and repair it.

Student: What about for the current limiting transistor?

SOPHIE: The base emitter voltage of Q8 is .349 volts.

The second student–SOPHIE interaction is a little less straightforward because, in addition to simulating the faulted power supply, SOPHIE must be able to parse the dialogue sufficiently to know that the student is asking for the base emitter voltage of transistor Q8. The meaning of the student's question depends heavily on the context established by his prior question.

Student: What should it be?
SOPHIE: In a working circuit, the base emitter voltage of Q8 is .593 volts.

The third student–SOPHIE interaction continues its dependence on SOPHIE's understanding of context and additionally requires SOPHIE to extract the answer by simulating both the faulted power supply and one that is operating correctly.

Student: Replace Q8.
SOPHIE: I am going to ask you some questions about how Q8 is faulted. Are any junctions shorted?

The fourth student–SOPHIE interaction indicates a clear step beyond what Brown et al. (1982) described as a knowledgeable system to what they considered to be an intelligent system. SOPHIE has progressed from a knowledgeable parsing of its dialogue with the student and simulation of various states of the power supply to a system exercising tutorial intelligence. It shadowed the student's solution path, modeled the student's troubleshooting hypotheses, determined that he is incorrect, elected to capture the dialogue initiative back from the student, and is undertaking a series of tutorial interactions intended to lead the student back to a more correct approach to the problem.

It is difficult to imagine any practical way to achieve this level of instructional functionality without the generative capability and mixed-initiative dialogue that distinguishes ITS from other forms of computer-based instruction.

3.2. What About Decision Aiding?

As can be seen from this example, an ITS such as SOPHIE has many elements and capabilities in common with a decision-aiding system, intended in this case to assist an electronics maintenance technician. The differences may lie entirely in the intentions of the user and depend on whether the system is being used to solve a problem or to effect a (more or less) persistent change in the user's knowledge and skills. The underlying computational capabilities appear to be the same.

Is it worth expanding the capabilities of an ITS to support decision aiding? Some evidence was provided by the development and evaluation by the U.S. Air Force of the Integrated Maintenance Information System (IMIS). This evidence was summarized by Teitelbaum and Orlansky (1996). They reviewed and analyzed the performance of technicians who were specially trained to troubleshoot and maintain F-16 avionics systems compared to the performance of less specifically (and much less expensively) trained general aircraft technicians.

The results of this comparison are shown in Table 1. It is notable that performance of the general aircraft technicians became essentially the same as that of the avionics specialists once they were given IMIS. Perhaps more notable, however, is that the performance of the general aircraft technicians using IMIS was superior to that of avionics specialists using the paper-based technical orders currently provided as maintenance aids. Finally, it may be worth noting that because IMIS was tied into the inventory and parts databases, it could generate the paperwork needed to order new parts in about a minute for both groups. IMIS provided both sophisticated decision aiding to the technicians as well as help in completing clerical tasks of a less exotic nature.

IMIS supported decision aiding for only three subsystems of F-16 avionics. Even at that, if its use were expanded to the full fleet of F-16s for these three subsystems, net cost savings to the Air Force would be about $23 million annually. There therefore appear to be substantial benefits in using ITS techniques and capabilities to provide both instruction and decision aiding.

The emphasis on IMIS as a decision-aiding, rather than just a maintenance-aiding, system is intended to suggest that this approach is not limited to maintenance or operation of devices. Its capabilities could be applied to more abstract undertakings, such as command and control in the military and management and administration in the civilian sector.

3.3. What's the Underlying Technical Idea?

Achieving these instructional and decision-aiding functionalities is an important goal, but it requires a technical idea to make it feasible. This idea was first articulated by Jaime Carbonell in 1970. He recommended replacing the ad hoc frame-oriented instructional approaches, such as those seen in programmed texts and best described as "intrinsic programming" (Crowder, 1959), with an information systems orientated approach. The ad hoc frame-oriented approach of intrinsic programming is widely used in technology-based instruction. It features a presentation (often a paragraph of text or a graphic of some sort) followed by questions. Answers by the student to these questions lead to different responses, usually different branching paths, by the system. The presentations, the questions, and all possible branching paths must be anticipated and prespecified by the instruction developers.

Carbonell's (1970) information systems orientated approach—or, as we might say today, an approach based on knowledge representation—allowed the instructional system to be generative. As suggested some time ago (Fletcher, 1975), the information systems orientated functionalities possessed by SOPHIE and other ITS involve computer representation

TABLE 1
Results From Trials of the Integrated Maintenance Information System

	Correct Solutions (%)		*Time to Solve Problems (min)*		*Time to Order Parts (min)*	
	Tech. Orders	*IMIS*	*Tech. Orders*	*IMIS*	*Tech. Orders*	*IMIS*
Avionics specialists	82	100	149	124	19	1
General aircraft tech.	69	98	176	124	25	1

Note. IMIS = Integrated Maintenance Information System.

or modeling of the student, the subject matter, and expert tutoring. This point of view now seems to be commonly accepted and is found to a significant degree in the architectures of ITS. Carbonell, Woolf and Regian (2000), and many others have emphasized knowledge representation as a key underlying capability for ITS. Progress is being made. Woolf and Regian emphasized that our ability to represent human cognition has gained considerable potency with advances made in cognitive science over the last 10 years. These advances should substantially enhance the promise and capabilities of ITS.

4. WHERE IS THIS TAKING US? ADL AND THE THIRD REVOLUTION IN LEARNING

The emphasis here on instructional technology brings us to revolutions in instruction. The first of these may have occurred with the development of written language about 7,000 years ago. It allowed the content of advanced ideas and teaching to transcend time and place. The second revolution in instruction began with the technology of books. Books made the content of high-quality instruction again available anywhere and anytime, but also inexpensive and thereby accessible to many more people. A third revolution in instruction appears to be accompanying the introduction of computer technology. The capability of this technology for real-time adjustment of instructional content, sequence, scope, difficulty, and style to meet the needs of individuals suggests a third pervasive and significant revolution in instruction. It makes both the content *and* the interactions of high-quality instruction widely and inexpensively accessible anytime, anywhere.

Building on this possibility, the ADL initiative is now being undertaken in the United States and elsewhere. This initiative is intended to provide education, training, and decision aiding (or “mentoring”) anytime, anywhere. It capitalizes on the growth of electronic commerce and the World Wide Web. It takes advantage of this global, almost irresistible activity, accelerates it, and applies it to learning. It will help ensure the availability of human competence in many organizations as people face the challenges of the 21st century.

The ADL initiative has been tasked to provide guidelines and specifications for all instructional activities of the U.S. Federal Government. It is similarly expected to provide guidelines and specifications for the instructional activities of a variety of countries in Europe and in the Pacific Rim. It may well succeed because it is not imposing standards developed by the government, but is pulling together the best ideas, guidelines, and specifications developed by a variety of industrial, academic, and government sources in the United States and elsewhere. It has become a massive cooperative development involving all economic sectors focused on achieving the goals of the ADL initiative, which are to provide learning and decision-aiding capabilities tailored to the needs of individuals anywhere and anytime they are needed.

5. WHAT DOES ADL HAVE TO DO WITH ITS?

“Intelligent” in the context of ITS refers as much to our intentions as our results. But it is more than a marketing term. It refers to the specific functionalities that are the goals of ITS development. As discussed, these functionalities are distinct from those found in more conventional approaches to computer-based instruction. They require ITS to generate instruction in real time and on demand as required by individual learners. This generative approach is also the goal of the ADL initiative, which is combining the benefits of both object-oriented development and Web delivery with those of technology-based instruction to achieve its objectives.

Despite their common goals, the ADL initiative and the development of intelligent instructional systems have been pursued in parallel but separately. The ADL initiative has focused on the specification of sharable instructional objects that can be retrieved from and then delivered over the Web anywhere and anytime they are needed. They can be used—and reused—to tailor instruction and decision-aiding presentations to the needs of individual learners. They support the core functional capability targeted by the ADL initiative of making the benefits of tutorial instruction and decision aiding universally and affordably available. The common interests shared by the developers of ITS and those pursuing the ADL initiative have been noted on both sides.

The ADL initiative is guided by a vision, roughly illustrated in Figure 1, of a future in which everyone has an electronic personal learning associate. This device will assemble learning or decision-aiding presentations on demand and in real time—any time, anywhere. The presentations will be exactly tailored to the needs, capabilities, intentions, and learning state of each individual or group (e.g., crew, team, or staff) of individuals. Communication with the device will be based on natural language dialogue initiated either by the device or by its users. The device will be small enough to be carried in a shirt pocket, or it will be wearable. It will be used by individuals learning by themselves, in groups, or in classrooms. It will, of course, be wireless.

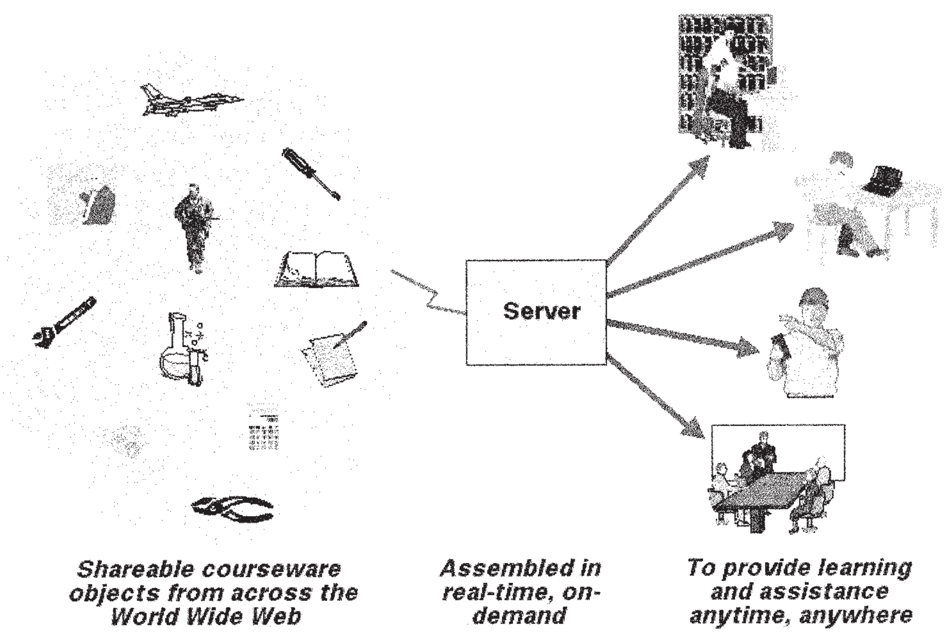

FIGURE 1 An ADL future.

Most of the technology needed to build such a device exists now. Although we cannot yet fit it into a shirt pocket, advances in electronics should take care of that. What is especially needed for instruction and decision aiding is content in the form of instructional objects, which we are calling *shareable instructional objects.* These objects, shown in the cloud on the left side of Figure 1, must be readily accessible across the World Wide Web—or whatever form our global information network takes in the future.

Once these objects exist, they must be identified, selected, and assembled in real time, on demand, and then handed to the personal learning associates, which provide the instruction or decision aiding. This work of identifying, selecting, and assembling objects is the job of the server, represented as a black box in the middle of Figure 1. By importing "logic" or instructional strategy objects, the server, which today would be called a learning management system, may acquire the capabilities of the intelligent tutoring or decision-aiding system we have been discussing. If it does so successfully, it will also provide the benefits we have been discussing. But in the ADL initiative, the underlying processes will rely on sharable instructional objects.

The ADL initiative and the development of ITS, then, have a number of key goals in common:

- Both are generative in that they envision the generation of presentations on demand, in real time.
- Both are intended to provide presentations that are tailored in content, sequence, level of difficulty, level of abstraction, style, and so on to users' intentions, backgrounds, and needs.
- Both have a stake in research intended to accomplish such individualization.
- Both can be used equally well to aid learning or decision making.
- Both are intended to accommodate mixed initiative dialogue in which either the technology or the user can initiate or respond to inquiries in natural language.
- Both will benefit greatly from a supply of sharable instructional objects readily available for the generation of instructional (or decision-aiding) presentations.

5.1. Sharable Instructional Objects

Roschelle and Kaput (1995) emphasized the promise of object-based software, such as that promoted by ADL, for combining many kinds of interactive content in multiple display formats and attaining for education the benefits now being realized in business from the use of integrated office software. Roschelle et al. (1999) illustrated these points by examining the software techniques underlying five already developed object-based education projects. Gibbons and Fairweather (2000) weighed these issues at length in examining the present status and future of computer-based instruction.

The cloud on the left side of Figure 1 represents the World Wide Web or whatever we use in the future to provide our global communication ether. One crucial matter for the implementation of ADL is what has loosely been called "content" in the form of shareable objects represented by the various icons shown in the cloud. People involved in ADL have spent—and continue to spend—much time, effort, and energy discussing what these sharable content objects (SCOs) should be. This matter transcends the immediate issues of ADL as evidenced by discussions edited by Wiley (2000) and others.

As presently defined, SCOs could be entire courses, lessons within courses, or modules within lessons. They could be electronic representations of media, text, images, sound, Web pages, or other data that can be presented to students. The size, or "granularity," of SCOs is a matter of considerable discussion. Gibbons, Nelson, and Richards (2000) emphasized that SCOs will be most useful if they are prepared in sufficiently small chunks to be accessed and reused by other instructional materials. Access to them must be rapid and easily accomplished across whatever form our global information network takes in the future. As suggested by Figure 1, once these SCOs exist, they must be available for assembly in real time and on demand by some sort of servers (middle of Figure 1) and then handed to client personnel learning associates (right side of Figure 1).

SCOs could also be material that is not seen by students, but is needed to register them for courses, report on their progress, collect them into classes and other administrative groupings, or store data needed to tailor instruction to individual student needs. Significantly they could also be content in the form of algorithms that aggregate, integrate, and sequence other objects as needed to manage the progress of students toward their attainment of specific instructional outcomes.

Some researchers have provided examples and suggested some architectures for this possibility. In their discussion of tutor agents, Koedinger, Suthers, and Forbus (1999) emphasized their value for "higher order" reasoning such as those associated with developing and assessing experimental strategies, determining representation strategies, developing conjectures, drawing conclusions, and making arguments in a science learning space. Ritter and Koedinger (1996) suggested a general architecture for agents that translate the mouse clicks, key presses, and other responses of students into semantic descriptions that can be used for tutoring and managing student progress toward curriculum goals.

These activities are the kinds of functions that are now expected to be performed by learning management systems, which are intended to be generalized capabilities for many different types of students, subject matters, curriculum structures, instructional approaches, and curriculum objectives. However, there seems to be a clear trade-off between the amount of functionality that can be built into these learning management systems and their generality. The more functionality built into them, the less general purpose they become. Learning management systems that import these functionalities as objects are likely to be more flexible, powerful, and generalizable.

In this sense, current discussion of content and what SCOs might be echoes the controversy that occurred early in the development of digital computers over whether data (traditional content) and logic (algorithms) should be stored in different ways and in different locations of the early machines (Goldstine, 1972). Goldstine reported that discussion was settled by John von Neumann, who recommended storing data and logic together, all as digital bits in a common memory. The nature of SCOs, their disposition, and whether to apply a solution analogous to von Neumann's all remain to be determined.

Still, SCOs have become the foundation for ADL. The availability of SCOs will significantly reduce the costs of preparing instruction, decision aiding, and job performance assistance for technology-based delivery. This is likely to be true whether the SCOs are assembled in advance by course authors and developers or, as suggested in Figure 1, they are assembled on demand and in real time by server algorithms incorporated in or imported into learning management systems. For these reasons, ADL development is presently focused on packaging SCOs in anticipation of what is called by Spohrer, Sumner, and Shum (1998), among others, the "educational object economy." The primary idea behind such an economy is that the emphasis in preparing materials for technology-based instruction (or decision aiding) will shift from the current concern with preparing content components, or instructional objects, to one of integrating already available content into meaningful and relevant presentations.

ADL developers recognize that this software engineering concern is only a beginning. The primary goal of ADL is not to promote tinkering with software objects but to develop the functional capability of producing instructional outcomes anytime, anywhere they are needed or desired. The ADL initiative has made substantial progress at this software engineering level, but it must also address real learning issues—it must determine how to assemble instructional objects to achieve targeted instructional objectives.

These two areas of software engineering and instructional design are not independent. They must be coordinated and "harmonized" to achieve the ADL vision. This point suggests that designers and developers need some understanding of the underlying ADL architecture and its software specifications to ensure that ADL presentations bring about efficient and effective learning. It also suggests that the software engineers designing SCOs along with their meta-data packaging and assumptions concerning course structures, learning management system structures, and communication protocols must understand and accommodate all varieties of instructional approaches and what each requires for software support.

This vision, which foresees the development of an instructional object economy, must start with the specification of sharable instructional objects. Fortunately, a great many individuals (technicians, software engineers, instruction designers, cognitive researchers, etc.) coming from organizations representing all sectors of the economy (government, industry, and academic) in the United States and elsewhere in Europe and Asia have joined in this quest. Development of these objects has become a global effort. Those involved in the United States ADL initiative mostly need to orchestrate the effort and document its results.

Shareable instructional objects that support ADL functional requirements must meet some criteria. Among these criteria, four seem most prominent:

- It must be possible to find needed and shareable objects. They must be *accessible.* Basically, we need widely accepted and standard ways to store objects so that widely accepted and standard ways can be used to find and retrieve them.
- Once found, the objects should be usable. This means that they must be *interoperable* and portable across most, if not all, platforms, operating systems, browsers, and courseware tools.
- Once implemented, the objects should continue to operate reliably. If the underlying platform, operating system, or browser is modified (for instance when a new version is released and installed), courseware objects should continue to operate as before. They should be *durable.*
- Finally, courseware objects should be *reusable.* Other platforms, operating systems, browsers, and courseware tools should be able to reuse, and perhaps even modify as needed, the original courseware objects.

Specifications intended to meet these criteria are being developed through an evolving series called the Sharable Content Objects Reference Model. Evolving versions of this model can be found at the ADL Web site (www.adlnet.org).

Specification of these objects arises from the intersection of disciplines concerned with learning and cognition, information storage and retrieval, and software engineering. How the objects should be packaged so that they are accessible depends on the model of cognition and learning or the learning strategy trying to find them. This packaging requires the development of a taxonomy that is sufficiently robust to accommodate a great many models and strategies. Aside from such models of human cognition, learning, and instruction, some representation of the physical world is also required to accommodate comprehensive range of objects. Finally, interoperability and reusability require techniques of software engineering that are compatible with the prevailing global ether—which today is the World Wide Web—as well as with the many operating environments in which the instructional objects will be used. Given the difficulty of the undertaking, it is fortunate that it has become a global quest. Current progress toward these ends can be viewed through many portals. The ADL Web site is as good a place to start as any.

6. SUMMARY

This discussion suggests that

- Substantial improvements in instructional effectiveness may be obtained by tailoring instruction to the needs and capabilities of individual learners. Evidence for these improvements may be found in studies of one-on-one tutoring. However, the provision of one instructor for each student is, in most applications, prohibitively expensive.
- The benefits of individualized instruction can be made affordable through the use of technology. Evidence of these benefits is provided by evaluations of both computer-based instruction in general and ITS in particular.
- Combined with the availability of sharable instructional objects available from the Web, ITS will lead to the development and wide use of personal learning associates, allowing high-quality instruction and decision aiding to become ubiquitous and affordable.
- This is a desirable result.

If Kurzweil's (1999) projections are correct, computers may become more effective in providing instruction than human tutors even if humans use all the techniques Graesser and his colleagues (1995) found they now neglect. We may not be implanting integrated circuits in our brains as Kurzweil suggested, but the goal of using technology to discover more than any human agent can about the unique potential of every individual and devising effective and individualized procedures to reach it seems both an appealing and realistic prospect.

REFERENCES

Atkinson, R. C., & Fletcher, J. D. (1972). Teaching children to read with a computer. *The Reading Teacher, 25,* 319–327.

Bloom, B. S. (1984). The 2-sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher, 13*(6), 4–16.

Brown, J. S., Burton, R. R., & DeKleer, J. (1982). Pedagogical, natural language and knowledge engineering in SOPHIE I, II, and III. In D. Sleeman & J. S. Brown (Eds.), *Intelligent tutoring systems* (pp. 227–282). New York: Academic.

Carbonell, J. R. (1970). AI in CAI: An artificial intelligence approach to computer-assisted instruction. *IEEE Transactions on Man–Machine Systems, 11,* 190–202.

Carroll, J. B. (1970). Problems of measurement related to the concept of learning for mastery. *Educational Horizons, 48,* 71–80.

Crowder, N. W. (1959). Automatic tutoring by means of intrinsic programming. In E. H. Galanter (Ed.), *Automatic teaching: The state of the art* (pp. 58–66). New York: Wiley.

Farr, M. J., & Psotka, J. (Eds.). (1992). *Intelligent instruction by computer: Theory and practice.* New York: Taylor & Francis.

Fletcher, J. D. (1975). Artificial intelligence applications in computer-based instruction. In A. G. Smith (Ed.), *Proceedings of ADCIS summer meeting* (pp. 43–62). Portland: University of Maine.

Fletcher, J. D. (1997). What have we learned about computer based instruction in military training? In R. J. Seidel & P. R. Chatelier (Eds.), *Virtual reality, training's future?* (pp. 169–177). New York: Plenum.

Fletcher, J. D., & Rockway, M. R. (1986). Computer-based training in the military. In J. A. Ellis (Ed.), *Military contributions to instructional technology* (pp. 171–222). New York: Praeger.

Gettinger, M., & White, M. A. (1980). Evaluating curriculum fit with class ability. *Journal of Educational Psychology, 72,* 338–344.

Gibbons, A. S., & Fairweather, P. G. (2000). Computer based instruction. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 410–442). New York: Macmillan.

Gibbons, A. S., Nelson, J., & Richards, R. (2000). The nature and origin of instructional objects. In D. Wiley (Ed.), *The instructional use of learning objects.* Retrieved March 2, 2001, from http//:www.reusability.org/read

Goldstine, H. H. (1972). *The computer from Pascal to von Neumann.* Princeton, NJ: Princeton University Press.

Gott, S. P., Kane, R. S., & Lesgold, A. (1995). *Tutoring for transfer of technical competence (AL/HR–TP–1995–0002).* Brooks AFB, TX: Armstrong Laboratory, Human Resources Directorate.

Graesser, A. C., & Person, N. K. (1994). Question asking during tutoring. *American Educational Research Journal, 31,* 104–137.

Graesser, A. C., Person, N. K., & Magliano, J. P. (1995). Collaborative dialogue patterns in naturalistic one-to-one tutoring. *Applied Cognitive Psychology, 9,* 1–28.

Heuston, D. H. (1997). *School improvement models: The manual model and the speed of light*. Sandy, UT: Waterford Institute.

Koedinger, K. R., Suthers, D. D., & Forbus, K. D. (1999). Component-based construction of a science learning space. *International Journal of Artificial Intelligence in Education, 10,* 292–313.

Kulik, J. A. (1994). Meta-analytic studies of findings on computer-based Instruction. In E. L. Baker & H. F. O'Neil, Jr. (Eds.), *Technology assessment in education and training* (pp. 9–34). Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Kurzweil, R. (1999). *The age of spiritual machines: When computers exceed human intelligence.* New York: Viking.

National Research Council. (1997). *Technology for the United States Navy and Marine Corps, 2000–2035: Vol. 4. Human Resources.* Washington, DC: Naval Studies Board, National Research Council, 1997.

Neisser, U. (1967). *Cognitive psychology.* New York: Appleton-Century-Crofts.

Orlansky, J., & String, J. (1979). *Cost-effectiveness of computer based instruction in military training* (IDA Paper P–1375). Alexandria, VA: Institute for Defense Analyses.

Psotka, J., Massey, L. D., & Mutter, S. A. (Eds.). (1988). *Intelligent tutoring systems: Lessons learned.* Hillsdale, NJ: Lawrence Erlbaum Associates, Inc.

Ritter, S., & Koedinger, K. R. (1996). An architecture for plug-in agents. *Journal of Artificial Intelligence in Education, 7,* 315–347.

Roschelle, J., DiGiano, C., Koutlis, M., Repenning, A., Phillips, J., Jackiw, N., & Suthers, D. (1999). Developing educational software components. *Computer, 32,* 2–10.

Roschelle, J., & Kaput, J. (1995). Educational software architecture and systemic impact: The promise of component software. *Journal of Educational Computing Research, 14,* 217–228.

Scriven, M. (1975). Problems and prospects for individualization. In H. Talmage (Ed.), *Systems of individualized education* (pp. 199–210). Berkeley, CA: McCutchan.

Sleeman, D., & Brown, J. S. (Eds.). (1982). *Intelligent tutoring systems.* New York: Academic.

Spohrer, J., Sumner, T., & Shum, S. B. (1998, October). Educational authoring tools and the educational object economy: Introduction to the special issue from the East/West Group. *Journal of Interactive Media in Education*. Retrieved March 2, 2001, from http://www-jime.open.ac.uk/98/10/spohrer-98-10-paper.html

Suppes, P. (1964). Modern learning theory and the elementary-school curriculum. *American Educational Research Journal, 1,* 79–93.

Suppes, P. (Ed.). (1981). *University-level computer assisted instruction at Stanford: 1968–1980.* Stanford, CA: Institute for Mathematical Studies in the Social Sciences.

Suppes, P., Fletcher, J. D., & Zanotti, M. (1975). Performance models of American Indian students on computer-assisted instruction in elementary mathematics. *Instructional Science, 4,* 303–313.

Suppes, P., Fletcher, J. D., & Zanotti, M. (1976). Models of individual trajectories in computer-assisted instruction for deaf students. *Journal of Educational Psychology, 68,* 117–127.

Suppes, P., & Morningstar, M. (1972). *Computer-assisted instruction at Stanford 1966–68: Data, models, and evaluation of the arithmetic programs.* New York: Academic.

Teitelbaum, D., & Orlansky, J. (1996). *Costs and benefits of the Integrated Maintenance Information System (IMIS) (IDA Paper P–3173).* Alexandria, VA: Institute for Defense Analyses.

Tobias, S. (1989). Another look at research on the adaptation of instruction to student characteristics. *Educational Psychologist, 24,* 213–227.

Weiner, N. (1954). *The human use of human beings.* Boston: Houghton Mifflin.

Wiley, D. (2000). *The instructional use of learning objects.* Retrieved February 12, 2001, from http://www.reusability.org/read

Woolf, B. P., & Regian, J. W. (2000). Knowledge-based training systems and the engineering of instruction. In S. Tobias & J. D. Fletcher (Eds.), *Training and retraining: A handbook for business, industry, government, and the military* (pp. 339–356). New York: Macmillan.